

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 30-12-2014		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 21-Aug-2012 - 20-Aug-2013	
4. TITLE AND SUBTITLE Final Report for "Hardware-Assisted Large-Scale Neuroevolution for Multiagent Learning"			5a. CONTRACT NUMBER W911NF-12-1-0439		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611103		
6. AUTHORS Mingjie Lin			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Central Florida 12201 Research Parkway, Suite 501 Orlando, FL 32826 -3246			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 61539-NS-RIP.1		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This DURIP equipment award was used to purchase, install, and bring on-line two Berkeley Emulation Engines (BEEs) and two mini-BEE machines to establish an FPGA-based high-performance multiagent training platform and its associated software. This acquisition of BEE4-W (Berkeley Emulation Engine) hardware platforms has costed \$201,500.00 in total. To accelerate both the multiagent software simulation and hardware development, a DELL integrated cluster environment unit was purchased, featuring 102.28 GHz processors (on 24 nodes) with 284 GB high-speed RAM.					
15. SUBJECT TERMS FPGA-based Multiagent Training Platform; Probabilistic Domain Transformation; Hardware-Assisted; FPGA; BEE; Hive Brain; Multiagent.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Mingjie Lin
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 407-882-2298

Report Title

Final Report for "Hardware-Assisted Large-Scale Neuroevolution for Multiagent Learning"

ABSTRACT

This DURIP equipment award was used to purchase, install, and bring on-line two Berkeley Emulation Engines (BEEs) and two mini-BEE machines to establish an FPGA-based high-performance multiagent training platform and its associated software. This acquisition of BEE4-W (Berkeley Emulation Engine) hardware platforms has costed \$201,500.00 in total.

To accelerate both the multiagent software simulation and hardware development, a DELL integrated cluster environment unit was purchased, featuring 192 2.8 GHz processors (on 24 nodes) with 384 GB high-speed RAM on an Infiniband backbone. We also installed cluster control software and licenses for the UBUNTU-based Linux operating system. Software was installed including JAVA and C/C++ compiler suites (debuggers, performance measurements, etc.), the Xilinx ISE software suite, and the Matlab software. The system was ordered in late 2012 and arrived in early 2013 and was up and running routinely with local users in a few weeks. This delay of equipment arrival is largely due to the fact that all

BeeCube systems are highly customized hardware. Numerous research groups at UCF have made extensive use of the DURIP cluster, running several different versions of Xilinx ISE design suite. This DURIP award has greatly increased turnaround time and productivity for these groups. Specifically, the purchased equipments have enhanced ongoing DARPA and ARO funded activities. Several exploratory research tasks using these new equipments are currently undergoing at UCF to 1) alleviate the computing performance bottleneck imposed by software-based simulation while conducting multiagent learning research, 2) significantly increase the multiagent HyperNEAT learning intensity to facilitate solving real-world problems, 3) enable training teams at much larger sizes than previously possible in simulation, and 4) make training adaptive neural networks with genuine synaptic plasticity feasible. In the long term,

this equipment acquisition can expand the research capability at the University of Central Florida (UCF) in the area of Robotics, Artificial Intelligence (AI), and evolvable hardware by forging a close collaboration between two research teams at UCF: one focusing on multiagent robot training (based on HyperNEAT technology and a hive brain) and the other on high-performance reconfigurable computing (BCM, MARC, and evolvable hardware).

To date, we have built the proposed hardware platform largely based on the existing technology of our Evolutionary Complexity Research Group (EPlex) and High-Performance Reconfigurable Computing Laboratory. Currently, we are focusing on developing a probabilistic computing engine for high-performance convolutions. This is motivated by the fact that both convolution and correlation are fundamental building blocks in scalable multiagent learning and training a multiagent hive brain. In addition, this equipment acquisition has enabled us for the first time to set up a practical educational platform at UCF. In this semester, we have experimented with using the purchased platforms for hardware software co-design class. In the long term, we believe that all these outcome from the proposed equipment acquisition can benefit various tasks relevant to DoD missions in the field, such as dynamic cordoning, coordinated surveillance, and multi-robot teleoperation.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Names of Post Doctorates

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Names of Faculty Supported

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Names of Under Graduate students supported

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PhDs

<u>NAME</u>

Total Number:

Names of other research staff

<u>NAME</u>

<u>PERCENT SUPPORTED</u>

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Research Enhanced:

- DARPA CSSG Phase 3: Real-world Scalable Multiagent Learning for Coordinated UGV Operations; Period: 8/11 - 8/12; Amount: \$249,347; PI: Kenneth O. Stanley; Program Manager: James Donlon.
- ARO Network Sciences Division: Training a Multiagent Hive Brain for Coordinated UGV Operations; Period: 9/11-9/14; Amount: \$150,000 in first year; PI: Kenneth O. Stanley; Program Manager: Purush Iyer

Technology Transfer

NA

ARO-DURIP Final Report:
Hardware-Assisted Large-Scale Neuroevolution
for Multiagent Learning
ARO Grant #: 61539-NS-RIP

Contents

1	Executive Summary	1
2	Research Enhanced	2
3	Description of Purchased Equipments	2
4	Examples of Usage	3
4.1	Main Task I: High-Throughput FPGA-based Multiagent Training Platform .	3
4.2	Main Task II: Energy-Efficient Multiplier-Less Discrete Convolver through Probabilistic Domain Transformation	5
5	Summary	5

1 Executive Summary

This DURIP equipment award was used to purchase, install, and bring on-line two Berkeley Emulation Engines (BEEs) and two mini-BEE machines to establish an FPGA-based high-performance multiagent training platform and its associated software. This acquisition of BEE4-W (Berkeley Emulation Engine) hardware platforms has costed \$201,500.00 in total. To accelerate both the multiagent software simulation and hardware development, a DELL integrated cluster environment unit was purchased, featuring 192 2.8 GHz processors (on 24 nodes) with 384 GB high-speed RAM on an Infiniband backbone. We also installed cluster control software and licenses for the UBUNTU-based Linux operating system. Software was installed including JAVA and C/C++ compiler suites (debuggers, performance measurements, etc.), the Xilinx ISE software suite, and the Matlab software. The system was ordered in late 2012 and arrived in early 2013 and was up and running routinely with local users in a few weeks. This delay of equipment arrival is largely due to the fact that all BeeCube systems are highly customized hardware. Numerous research groups at UCF have made extensive use of the DURIP cluster, running several different versions of Xilinx ISE design suite. This DURIP award has greatly increased turnaround time and productivity for these groups.

Specifically, the purchased equipments have enhanced ongoing DARPA and ARO funded activities. Several exploratory research tasks using these new equipments are currently undergoing at UCF to 1) alleviate the computing performance bottleneck imposed by software-based simulation while conducting multiagent learning research, 2) significantly increase the multiagent HyperNEAT learning intensity to facilitate solving real-world problems, 3) enable training teams at much larger sizes than previously possible in simulation, and 4) make training adaptive neural networks with genuine synaptic plasticity feasible. In the long term, this equipment acquisition can expand the research capability at the University of Central Florida (UCF) in the area of Robotics, Artificial Intelligence (AI), and evolvable hardware by forging a close collaboration between two research teams at UCF: one focusing on multiagent robot training (based on HyperNEAT technology and a hive brain) and the other on high-performance reconfigurable computing (BCM, MARC, and evolvable hardware).

To date, we have built the proposed hardware platform largely based on the existing technology of our Evolutionary Complexity Research Group (EPlex) and High-Performance Reconfigurable Computing Laboratory. Currently, we are focusing on developing a probabilistic computing engine for high-performance convolutions. This is motivated by the fact that both convolution and correlation are fundamental building blocks in scalable multiagent learning and training a multiagent hive brain. In addition, this equipment acquisition has enabled us for the first time to set up a practical educational platform at UCF. In this semester, we have experimented with using the purchased platforms for hardware software co-design class. In the long term, we believe that all these outcome from the proposed equipment acquisition can benefit various tasks relevant to DoD missions in the field, such as dynamic cordoning, coordinated surveillance, and multi-robot teleoperation.

Key Words: Hardware-Assisted; FPGA; BEE; Hive Brain; Multiagent.

2 Research Enhanced

- DARPA CSSG Phase 3: *Real-world Scalable Multiagent Learning for Coordinated UGV Operations*; Period: 8/11 - 8/12; Amount: \$249,347; PI: Kenneth O. Stanley; Program Manager: James Donlon.
- ARO Network Sciences Division: *Training a Multiagent Hive Brain for Coordinated UGV Operations*; Period: 9/11-9/14; Amount: \$150,000 in first year; PI: Kenneth O. Stanley; Program Manager: Purush Iyer

3 Description of Purchased Equipments

We have purchased the following equipments as proposed in our DURIP project. Table 1 lists their short descriptions and price information.

No.	Item #	Description	Qty	Unit Price	Total
1	BEE4-LX550T-2C-CM	BEE4 LX550T Configuration for University Only	2	\$ 100,750	\$ 201,500
2	BEE4-DRAM-64G-CM	64GB DDR3-1066 x4 RDIMMs for BEE4 system		included	
3	BEE4-PCIE-LP-CM	PCI Express Kit w/ low-profile bracket		included	
4	BEE4-WTY-1Y-CM	1 year limited hardware warrant		included	
				Total	\$201,500.00

Table 1: Equipment list.

BEE (Berkeley Emulation Engine) Platform

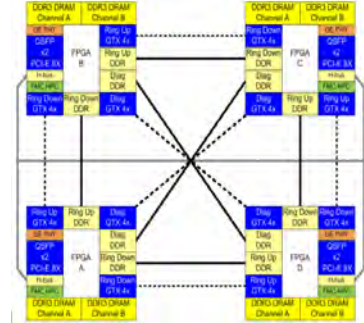


Figure 1: Interconnect Network Architecture of a BEE platform.

We purchased the Berkeley Emulation Engine (BEE 4) platforms to establish an FPGA-based high-performance multiagent training platform. The Berkeley Emulation Engine (BEE4) is a commercial, stackable full speed multi-FPGA based prototyping platform, integrated with DAC/ADC modules for mixed signal and digital communications designs.

Specifically, each BEE4 system or module consists of 4 Xilinx Virtex-6 FPGAs based on a unique HPC (High Performance Computing) Honeycomb architecture, which allow us to build application-specific computing platform to replace a large number of computers. Each BEE4 module supports a variety of high-end Virtex-6 FPGAs, including LXT 240/365/550 and SXT 315/475, allowing BEE4 to support 20 MGate designs per module, or 400 MGates per rack, all of which are interconnected through high-speed links. This FPGA-based computing platform can 1) *significantly increase the number of evaluations possible in multiagent HyperNEAT learning to facilitate solving real-world problems*, 2) *enable training teams at much larger sizes than previously possible in simulation*, and 3) *make adaptive neural networks through synaptic plasticity feasible*. All these benefits can alleviate the computing performance bottleneck imposed by software-based simulation while conducting multiagent learning research. In addition, BEE4 modules can be stacked or clustered to increase capacity without losing speed, which is critical as we scale up the size of our multiagent robot training system. Finally, because of the versatility of the BEE4 platform, this facility can serve the research and teaching needs of several faculty members in our department and enhance our research in these areas.

4 Examples of Usage

The two DARPA-sponsored projects, the *Real-world Scalable Multiagent Learning for Coordinated UGV Operations* (Period: 8/11 - 8/12; Amount: \$249,347; PI: Kenneth O. Stanley; Program Manager: James Donlon) and the *Training a Multiagent Hive Brain for Coordinated UGV Operations* (Period: 9/11-9/14; Amount: \$150,000 in first year; PI: Kenneth O. Stanley; Program Manager: Purush Iyer) have made extensive use of the computational cluster. (These projects are given highest priority usage.) When the hardware emulation engines constructed with BEE4 platforms are finished, these two projects will use the hardware platforms extensively.

4.1 Main Task I: High-Throughput FPGA-based Multiagent Training Platform

We are currently using the reconfigurable hardware within the purchased BEE4 platforms to construct a high throughput FPGA-based multiagent training platform capable of evaluating probabilistic networks with arbitrary DAG (directed acyclic graph) topology. We have finished building the application-specific processor cores that are specifically tailored for multiagent training. We are working on constructing application-specific memory access network to significantly improve the memory bandwidth by avoiding “Von-Neumann bottleneck”. Eventually, we plan to replace networks between clustered computer with high-speed inter-FPGA links with 200 times more bandwidth.

The above scheme has been built on the top of our previous work of building large-scale Bayesian Computing Machine using reconfigurable logic, where we successfully

developed a high throughput Bayesian computing machine (BCM) that not only can be applied to many important algorithms in artificial intelligence, signal processing, and digital communications, but also has high reusability, i.e., a new application needs not change a BCM’s hardware design, only new task graph processing and code compilation are necessary. In this DURIP effort, we are currently working on make the computing architecture of BCM to scale linearly with the size of the FPGA on which it is implemented. Because of our success with Bayesian computing machine containing 16 processing nodes in a Virtex-5 FPGA (XCV5LX155T-2), especially because of the strong similarity in computing pattern between evolving neural network and Bayesian network computing, we are quite optimistic that the similar two orders of magnitude of performance improvement in BCM can also be achieved in the proposed multiagent training platform against a 2.4 GHz Core 2 Duo Intel processor and a GeForce 9400m using the CUDA software package. In addition, we expect that all these huge perform improvements are obtained with only tens of Watts while the reference GPU platform consumed more than 200 Watts.

The architecture of our hardware-assisted multiagent platform was highly optimized for training large-scale neural networks. More specifically:

- Each processing node only implements the minimum hardware required for the necessary neural network computation. Unlike in conventional parallel processors implemented with ASIC technology, where the pipeline stages seldom exceed 10 due to data dependencies, in the proposed multiagent training platform, we pipeline each processing node to 20~30 in order to maximize the throughput. Furthermore, instead of building sophisticated control circuitry to handle data or instruction hazards, the multiagent training platform completely avoids dependencies by pre-processing the computing task graph and strategically allocating memory for intermediate results during the execution.
- The memory of the multiagent neural substrates is physically distributed in order to take advantage of the massively distributed memory blocks in a modern FPGA. The communication between processors and memories is handled by two separate crossbar switches. Combined with pre-processing of the computing task graph in the compilation phase, our novel memory allocation scheme can effectively avoid any data dependency and memory hazards.
- As for each processing node, the scheduler that controls the two crossbars is also deeply pipelined to improve the throughput. Given the fact that all information needed for memory allocation of outgoing messages is available before its results, we can start computing the schedule right after the first stage of the processor.

4.2 Main Task II: Energy-Efficient Multiplier-Less Discrete Convolver through Probabilistic Domain Transformation

Large-scale convolution is a fundamental building block for all of the multiagent learning algorithms and many important video and image processing algorithms. Using this purchased BEE4 platform, we have designed a high-performance reconfigurable discrete convolver specifically designed for FPGA-based multiagent training processors. While the conventional multiplier-based architecture can only achieve $\mathcal{O}(N^2)$, our proposed architecture can achieve approximately $\mathcal{O}(N)$ in algorithmic complexity, therefore highly scalable and energy efficient. This significant reduction in algorithm complexity is made possible by using probabilistic domain transformation and several novel reconfigurable hardware techniques to convert computationally intensive multiplications into additions that manipulate random signal samples. In addition, these techniques make the proposed architecture highly fault-tolerant because information to be processed is encoded with probability density function instead of its binary forms. As such, the local perturbations of its computing accuracy or signal values are inconsequential to its overall results.

To validate this proposed probabilistic architecture for discrete convolver, We have developed several prototypes with Virtex 6 FPGA devices (XC6VLX550t). To further optimize the performance of our probabilistic convolver, we developed four novel hardware techniques—Segmented Memory Swapping, Stochastic Mixing Scheme, Virtual Indexing Scheme, and Scalable Adder Extension—to minimize hardware usage and mitigate memory bandwidth bottleneck. Our prototype of probabilistic convolver requires just $4.09\ \mu\text{s}$ to perform a 128×128 convolution and dissipates only 166.63 nJ in dynamic energy consumption at 250 MHz. We believe that this new architecture can be exploited in all the real-time applications in which energy-efficient convolutions are required and it can be realized with many other FPGA device families.

5 Summary

The DURIP project officially started on August 15, 2012. All equipments were purchased within one month. However, because all FPGA platforms are custom-built, the final arrival of all these BEE4 platforms was delayed to May 2013. Fortunately, the FPGA-based high-performance multiagent training platform will be extensively used for at least five years. After that, the platform will be upgraded.

The purchased platforms are housed at the Computer Architecture Laboratory (Harris Engineering Center Room 242 with additional resources in Room 238) in the School of Electrical Engineering and Computer Science at the University of Central Florida. The MARC training of defense-related engineering students has taken place through the lab sessions of course EEL5722c (Advance FPGA Design and Technology).

We have included the following as “Synergistic Benefits to DoD”:

- Research

- We developed the FPGA-based high-performance computing systems that alleviate the computing performance bottleneck imposed by software-based simulation while conducting multiagent learning research.
 - This system can significantly increase the multiagent HyperNEAT learning intensity to facilitate solving real-world problems.
 - Through this purchase, we have explored how to enable training teams at much larger sizes than previously possible in simulation, and make training adaptive neural networks with genuine synaptic plasticity feasible.
- Education (Both Graduate and Undergraduate)
 - For the first time, we set up a practical educational platform at UCF for many of our students who are currently working (or will work) for defense-related industry.
 - This DURIP project provides valuable opportunities for both graduate and undergraduate students to learn real-world hardware design.

All these outcome from the proposed equipment acquisition can benefit various tasks relevant to DoD missions in the field, such as dynamic cordoning, coordinated surveillance, and multi-robot teleoperation. Some of these benefits have already been realized.



Purchase Order

University of Central Florida

Orlando Tech Center
12479 Research Parkway
Orlando FL 32826
United States

Vendor: 0000071428
BEEcube Inc
39465 Paseo Padre Pkwy Suite 3700
Fremont CA 94538

DUPLICATE

Dispatch via Print

Purchase Order	Date	Revision	Page
UCF01-0000270931	06/11/2013		1
Payment Terms	Freight Terms	Ship Via	
Net 30	Destination, Freight Prepay	Best Way	
Buyer	Phone	Currency	
Brian Michael Sargent	407/823-3304	USD	

Ship To: See Detail Below

Bill To: 12424 Research Parkway, Suite 300
Orlando FL 32826-3249
United States

Tax Exempt? Y **Tax Exempt ID:** 85-8012703010C-9

Replenishment Option: Standard

Line-Sch	Item/Description	Mfg ID	Quantity	UOM	PO Price	Extended Amt	Due Date
1- 1	miniBEE Academic Edition (LX550T FPGA, 4GB DDR3) with a single FPGA chip donated by Xilinx University Program with a commercial value of \$4537.00. p/n mBEE-LX550T-U. Attn: Kenneth R Enloe Ship To: 4000 Central Florida Blvd Bldg 116, Room 346 Orlando FL 32816-2362 United States		2.00	EA	15,500.00	31,000.00	06/11/2013

Schedule Total 31,000.00

Item Total 31,000.00

2- 1	High performance workstation for miniBEE design environment, p/n HPC-WORKSTATION-U. Attn: Kenneth R Enloe Ship To: 4000 Central Florida Blvd Bldg 116, Room 346 Orlando FL 32816-2362 United States		2.00	EA	5,000.00	10,000.00	06/11/2013
------	--	--	------	----	----------	-----------	------------

Schedule Total 10,000.00

Name of requestor: Ken Enloe
Phone Number: 407-823-5790
Building & Room No.: HEC 346B
Dept. Name: EECS
FAX No: 407-823-5835
Project: 16226056

Requested by: Dr. Mingjie Lin, 407-882-2298, mingjie.lin@gmail.com

BEEcube Quote #UCF20130508 is attached.
Waiver of Competition attached.

Item Total 10,000.00

Total PO Amount 41,000.00

This P.O. is expressly made subject to the terms found at: <http://www.purchasing.ucf.edu> or call (407) 823-2661 for a copy. Any and all vendor terms are invalid unless approved in writing by UCF General Counsel.

Authorized Signature



Purchase Order

Page: 1 of 1

University of Central Florida

Orlando Tech Center
12479 Research Parkway
Orlando FL 32826
United States

Dispatch Via Email		
Purchase Order UCF01-0000267620	Date 2013-04-08	Revision
Payment Terms Net 30	Freight Terms Destination, Freight Prepay	Ship Via Best Way
Buyer Tangel, Christine Marie	Phone 407/823-5347	Currency USD

Vendor: 0000071428
BEEcube Inc
39465 Paseo Padre Pkwy
Suite 3700
Fremont CA 94538
United States

Ship To: See Detail Below

Attention: Kenneth R Enloe

Bill To: 12424 Research
Parkway, Suite 300
Orlando FL 32826-3249
United States

Tax Exempt? Y

Tax Exempt ID: 85-8012703010C-9

Replenishment Option: Standard

Line- Sch	Item/Description	Mfg ID	Quantity	UOM	PO Price	Extended Amt	Due Date
1 - 1	Berkeley Emulation Engine (BEE4), p/n BEE4-LX550T-2C-U, with accessories per quote.		2.00	EA	64869.00	129738.00	04/08/2013

Ship To: 4000 Central Florida
Blvd
Bldg 116, Room 346
Orlando FL 32816-
2362
United States

Schedule Total 129738.00

Name of requestor: Ken Enloe
Phone Number: 407-823-5790
Building & Room No.: HEC 346B
Dept. Name: EECS
FAX No: 407-823-5835
Project: 16226056

Requested by: Dr. Mingjie Lin, 407-882-2298, mingjie.lin@gmail.com

BEEcube Quote #UCF20130213 is attached.
Sole Source Certificate & Terms and Conditions attached.

Item Total 129738.00

Total PO Amount 129738.00

This P.O. is expressly made subject to the terms found at:
<http://www.purchasing.ucf.edu> or call (407) 823-2661 for a copy. Any and all
vendor terms are invalid unless approved in writing by UCF General Counsel.

Authorized Signature